

NO-A170 050

LEARNING THE HIDDEN STRUCTURE OF SPEECH(U) CALIFORNIA
UNIV SAN DIEGO LA JOLLA INST FOR COGNITIVE SCIENCE
J L ELMAN ET AL. FEB 87 ICS-8701 N00014-85-K-0076


1/1

UNCLASSIFIED

F/G 5/7

ML

		4											
							0000						
							IND						
							q 8/						
							DIR						



XEROCOPY RESOLUTION TEST CHART

AD-A178 058

12

LEARNING THE HIDDEN
STRUCTURE OF SPEECH

J. L. Elman and D. Zipser

February 1987

ICS Report 8701

COGNITIVE
SCIENCE



UCSD

DTIC
ELECTE
MAR 20 1987
S A

This document has been approved
for public release and sale; its
distribution is unlimited.

INSTITUTE FOR COGNITIVE SCIENCE

UNIVERSITY OF CALIFORNIA, SAN DIEGO

LA JOLLA, CALIFORNIA 92093

DTIC FILE COPY

87 3 20 010

(12)

LEARNING THE HIDDEN STRUCTURE OF SPEECH

J. L. Elman and D. Zipser

February 1987

ICS Report 8701

*Jeffrey L. Elman
Department of Linguistics
University of California, San Diego
La Jolla, California 92093*

*David Zipser
Institute for Cognitive Science
University of California, San Diego
La Jolla, California 92093*

Authors' names are in alphabetical order. We wish to thank Dave Rumelhart for many useful discussions, and Jeffrey Teeters, Arthur Shapiro, and David Fare for programming assistance.

This work was supported by contract N00014-85-K-0076 from the Office of Naval Research, a grant from the System Development Foundation, and the Air Force Office of Scientific Research, Air Force Systems Command, USAF, grant number AFOSR-86-0062. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the sponsoring agencies. Approved for public release; distribution unlimited. Reproduction in whole or in part is permitted for any purpose of the United States Government. Requests for reprints should be sent to the Institute for Cognitive Science, C-015; University of California, San Diego; La Jolla, CA 92093. Copyright © 1986 by Jeffrey L. Elman and David Zipser.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified		1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; distribution unlimited	
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE			
4. PERFORMING ORGANIZATION REPORT NUMBER(S) ICS 8701		5. MONITORING ORGANIZATION REPORT NUMBER(S)	
6a. NAME OF PERFORMING ORGANIZATION The Regents of the University of California	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Personnel & Training Research Program ONR (Code 1142PT), 800 N. Quincy St.	
6c. ADDRESS (City, State, and ZIP Code) University of California, San Diego La Jolla, California 92093		7b. ADDRESS (City, State, and ZIP Code) Arlington, Virginia 22217-5000	
8a. NAME OF FUNDING / SPONSORING ORGANIZATION	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-85-K-0076	
8c. ADDRESS (City, State, and ZIP Code)		10. SOURCE OF FUNDING NUMBERS	
		PROGRAM ELEMENT NO 61153N	PROJECT NO. RR04206
		TASK NO 0A	WORK UNIT ACCESSION NO 667-542
11. TITLE (Include Security Classification) Learning the Hidden Structure of Speech			
12. PERSONAL AUTHOR(S) ELMAN, Jeffrey Locke; ZIPSER, David			
13a. TYPE OF REPORT Interim	13b. TIME COVERED FROM 12/1/85 TO 1/30/87	14. DATE OF REPORT (Year, Month, Day) 1987 February	15. PAGE COUNT 18
16. SUPPLEMENTARY NOTATION			
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	
05	10		
		parallel distributed processing, speech recognition, learning, back propagation, neural network	
19. ABSTRACT (Continue on reverse if necessary and identify by block number)			
<p>In the work described here, we apply the back-propagation neural network learning procedure to the analysis and recognition of speech. Because this learning procedure requires only examples of input-output pairs, it is not necessary to provide it with any initial description of speech features. Rather, the network develops its own set of representational features during the course of learning. A series of computer simulation studies were carried out to assess the ability of these networks to accurately label sounds; to learn to recognize sounds without labels; and to learn feature representations of continuous speech. These studies demonstrated that the networks can learn to label pre-segmented naive sounds tokens with accuracies of up to 95%. Networks trained on segmented sounds using a strategy that requires no external labels were able to recognize and delineate sounds in continuous speech. These networks developed rich internal representations that included units which corresponded to such traditional distinctions as vowels and consonants, as well as units which were sensitive to novel and</p> <p style="text-align: right;">(cont.)</p>			
20. DISTRIBUTION / AVAILABILITY OF ABSTRACT <input type="checkbox"/> UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION Unclassified	
22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. Harold Hawkins		22b. TELEPHONE (Include Area Code) (202) 696-4323	22c. OFFICE SYMBOL ONR 1142PT

19. (cont.)

non-standard features. Networks trained on a large corpus of unsegmented, continuous speech without labels also developed interesting feature representations that may be useful in both segmentation and label learning. The results of these studies, while preliminary, demonstrate that back-propagation learning can be used with complex, natural data to identify a feature structure that can serve as the basis for both analysis and non-trivial pattern recognition. These findings can make a significant contribution to the long-standing problems of speech feature decomposition and computer-based speech recognition.

Accession For

NTIS	<input checked="" type="checkbox"/>
DOE	<input type="checkbox"/>
Unpublished	<input type="checkbox"/>
Other	<input type="checkbox"/>

1/1



Learning the Hidden Structure of Speech

JEFFREY L. ELMAN and DAVID ZIPSER

INTRODUCTION

The recognition of speech is one of many things that is carried out by humans with apparent ease, but that has been done by computers only at great cost, with high error, and in highly constrained situations. Whether or not one is interested in machine-based speech recognition per se, the difficulties encountered by such systems may be diagnostic of flaws in the theoretical frameworks that motivate them. We believe that part of the difficulty in such systems lies in the use of inappropriate features as units for recognizing and representing speech. But what are the appropriate units and how are they to be found? In this paper we describe studies designed to determine whether these units can be learned. We use a newly developed learning procedure for artificial neural networks. This approach not only provides a way to learn to recognize speech, but also enables the required computations to be carried out in parallel in a brain-like fashion.

The question "What are the units of speech perception?" has been long-standing and controversial. The problem is that when one looks at the acoustic waveform, there are rarely obvious clues as to the boundaries between segments, let alone an indication of what those segments are. The speech sound wave varies continuously and smoothly over time. There is nothing mysterious or unexpected in this; it is the acoustic consequence of the fact that the production of speech involves a high degree of coarticulation (i.e., the mutual influence of neighboring sounds) with smooth transitions from one sound to another.

While one has the impression that speech is made up of concatenated "sounds," just what those sounds are is debatable. At least eight different levels of representation have been proposed to intervene between the speech wave and the representation "word": spectral templates (Lowerre, 1976), features (Cole, Stern, & Lasry, 1986), diphones (Dixon & Silverman, 1976; Klatt, 1980), context-sensitive allophones (Wickelgren, 1969), phonemes (Pisoni, 1981), demisyllables (Fujimura & Lovins, 1978), syllables (Mehler, 1981), and morphemes (Aronoff, 1976; Klatt, 1980). While these are all reasonable candidates for representing speech, the problem is that they have not been derived, in a canonical way, from the speech data itself. Learning provides a systematic way to find recognition features in data.

Whether or not the representations used in the perception of speech are innate or learned remains open, and we do not wish to take a strong position on this issue. However, much of the motivation for supposing that internal representations in perception are innate has come from the apparent poverty of data and the weakness of learning algorithms. Recent developments in parallel distributed processing (PDP) learning algorithms have demonstrated that a surprisingly small amount of data may contain sufficient cues to its intrinsic structure so that rather simple learning rules can be used to induce this

structure. In the current paper we attempt to demonstrate the consequences of requiring that representations be learnable. To do this we have taught trainable networks a series of speech recognition tasks and then examined the internal representations that are generated. The networks are extremely adept at solving the recognition tasks. They spontaneously develop their own representations which sometimes, but not always, correspond to our previous categories of representational units.

BACK PROPAGATION OF ERROR

In these studies we use a form of the "generalized delta-rule" known as "back propagation of error" (Rumelhart, Hinton, & Williams, 1986). This algorithm provides a technique for training a multilayer network to associate many pairs of patterns. The complete set of pattern pairs to be learned can be thought of as the extensional definition of a vector-valued function whose domain is the set of input patterns and whose range is the set of outputs.

As learning schemes become more powerful, the set of functions that can be learned increases. For example, the perceptron convergence procedure (Rosenblatt, 1962) can program networks to compute linear Boolean functions such as AND and OR but not nonlinear ones such as XOR. The early generalizations of the perceptron rule that extended the learning set to patterns with continuous rather than Boolean values are also limited to learning linear functions. Back propagation, on the other hand, can program multilayer networks to compute all the Boolean functions. Since it is applicable to patterns with continuous component values, back propagation can also deal with a much wider range of functions.

Much of the significance of back-propagation learning stems from the fact that it is defined on a neural-like network. An example of such a network is shown in Figure 1. The output of each unit is a function of the weighted sum of its inputs. It is these weights that are changed as learning proceeds. Each layer in the back-propagation network can have any number of units. In the work described here only strictly layered networks are used in which each unit in a layer receives inputs from every unit in the layer below. Training proceeds in cycles. In each cycle a pattern pair is chosen from the function definition. The input pattern is applied to the first layer of the network, and the activity it generates is passed successively to the other layers until it produces a pattern of activity on the output layer. This output pattern vector is then subtracted from the correct output pattern to produce an error pattern. This error in turn is used to adjust the weights in the output layer and then (by back propagation) to adjust the weights in lower layers. The exact way the error is propagated back down through each layer constitutes the novel part of the learning rule. Discovering just how to do this correctly was the significant accomplishment of Rumelhart, Hinton, and Williams (1986). The details of their procedure, which we have used in this work, are given in the caption of Figure 1.

When computations are programmed inductively in this manner, values from the range of the function must be available. Frequently these values are supplied by an external source or "teacher," which assumes that such a teaching input is available; however, in certain situations, this assumption may be an unrealistic. This raises the question of how to configure a system so it can learn, either without an external teacher or with the kind of information more realistically available. Several solutions have been proposed. One of the simplest and most elegant of these is to use teaching patterns that are the same as the input or some fixed transformation of the input. While this would seem to limit us to learning the identity function (or some fixed transformation of it), it has been shown that with this procedure the hidden units learn to represent the input patterns in terms of salient features. When the number of hidden units is less than the number of input units, the information in the input is represented at a lower dimensionality. In many perceptual problems this lower dimensional feature representation is just what is needed as a basis for further processing. We use this approach in some of the speech recognition studies reported here.

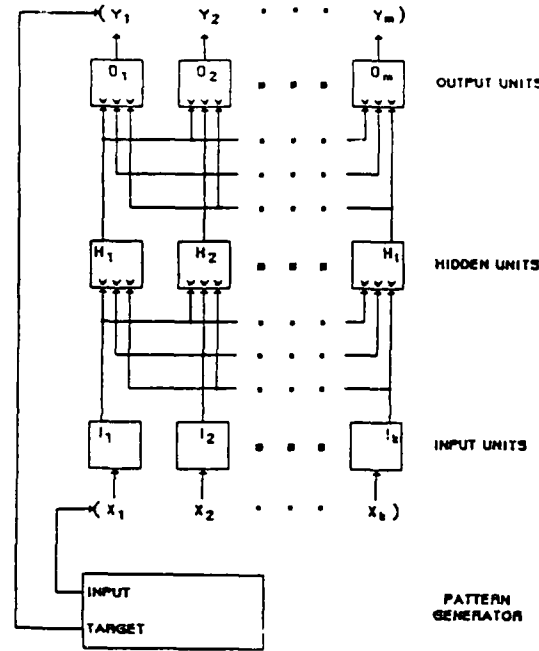


FIGURE 1. Strictly layered back propagation networks of the kind depicted in this figure were used for all the work described in this paper. Before training begins, all the weights and biases are set to random values between -1.0 and 1.0 . At the beginning of each learning cycle, i.e., at time $t = 0$ for that cycle, the pattern generator provides an input pattern (x_1, x_2, \dots, x_k) and a target pattern (y_1, y_2, \dots, y_m) . At time $t = 1$ the outputs of the input units are set equal the input pattern, that is: $I_1 = x_1, I_2 = x_2, \dots, I_k = x_k$. The output activities of the hidden units are set at time $t = 2$ to

$$H_i(t+1) = \text{squash}(\text{bias}_{H_i} + \sum_{j=1}^k W_{H_i, j} I_j(t))$$

where

$$\text{squash}(x) = \frac{1}{1 + e^{-x}}.$$

Similarly, the activity of the output units are set at $t = 3$ to

$$O_i(t+1) = \text{squash}(\text{bias}_{O_i} + \sum_{j=1}^l W_{O_i, j} H_j(t)).$$

All these values are held constant until the end of the current learning cycle. The next step is to change all the weights and biases according to the back propagation rule. At $t = 4$ the output unit weights are changed by the rule

$$\Delta W_{O_i, j} = \eta \delta_{O_i} H_j$$

$$\Delta \text{bias}_{O_i} = \eta \delta_{O_i}$$

$$\delta_{O_i} = (y_i - O_i) O_i (1 - O_i)$$

η = learning rate constant.

At $t = 5$ the hidden unit weights are changed by

$$\Delta W_{H_i, j} = \eta \delta_{H_i} I_j$$

$$\Delta \text{bias}_{H_i} = \eta \delta_{H_i}$$

$$\delta_{H_i} = \eta H_i (1 - H_i) \sum_{k=1}^m \delta_{O_k} W_{O_k, i}.$$

For the derivation of these rules see Rumelhart, Hinton, and Williams (1986).

PRELIMINARY CONSIDERATIONS

Before back propagation could be used to recognize speech it was necessary to find a way to present the sound to the network. The speech signal itself changes with time but the networks require fixed input and target samples. The approach we used here was to present the learning network with fixed input patterns each of which consists of a set of sequential samples. In some cases the individual samples in the input pattern were Fourier amplitude spectra; in other cases the actual digitalized sound samples were used. Preliminary studies were carried out to find the workable ranges for the number of frequencies and the amount of time that had to be represented in the input samples.

Another consideration was the way to normalize this data. The dynamic range of the Fourier spectra is large, and the vast majority of the points have very low values. We found that certain versions of the learning procedures had difficulty with this type of input, consisting of a vast sea of near zero values with a few high peaks. The situation was further complicated by large amplitude differences between examples of the same sound.

We used two different strategies in preprocessing the input data. One strategy involved finding ad hoc methods to deal with the problems raised by the amplitude and dynamic range. The other strategy was to train a network to do the preprocessing itself. Details of these strategies will be given later.

DIRECT LEARNING OF PHONETIC LABELS

In our first series of studies we asked how a back-propagation network might solve the problem of learning to label a set of highly confusable syllables. The basic idea was to use a spectrogram of a sound as the input pattern and a target bit pattern with one bit position for each of the types of sound. The task of the network was to learn to set the output unit corresponding to the sound type of the input to 1, while setting all the other output units to zero. We chose the syllables [ba], [bi], [bu], [da], [di], [du], [ga], [gi], and [gu] because this set of three voiced stops, paired with each of three vowels, is known to exhibit a high degree of variability due to coarticulation. Although listeners readily report all versions of (for example) the [d] as sounding the same, the acoustic patterns corresponding to the consonant differ greatly across the three vowel contexts. Indeed, this represents the paradigm case of perceptual invariance coupled with acoustic variability.

The stimuli for the experiment were prepared as follows. A single male speaker recorded a set of 505 tokens of the set of nine syllables (about 56 tokens of each syllable). Tokens were recorded in a moderately quiet environment, but with no particular effort at eliminating background noise; nor was an attempt made to ensure a constant rate of speech or uniformity of pronunciation. Recording was carried out through analog-to-digital conversion at a 10 kHz sampling rate and low-pass filtered at 3.5 kHz.

The beginning of each consonant was located and an FFT analysis was carried out over 6.4 ms frames, advancing 3.2 ms per frame for 20 frames. The output of the FFT was reformatted to give spectral magnitudes over 16 frequency ranges. As a result, each token was represented as a 16 x 20 array of positive values. Finally, these FFT magnitudes were normalized to the token average and "squashed" using the logistic. Examples of the resulting values for several tokens are graphed in Figure 2.

The network used consisted of 320 input units, between two and six hidden units (in different conditions), and as many output units as there were types to be labeled. The network used here and throughout this work was strictly layered; i.e., every unit of the input and hidden layer was connected to every unit on the layer above.

The data set of sounds was divided into two equal parts. One was used for training and the other was kept for testing performance on untrained examples. On each cycle of the teaching phase the following occurred: (a) One of the syllables from the training set was selected at random and applied to the input layer; (b) activation was propagated up from the input layer to the hidden layer, and from the

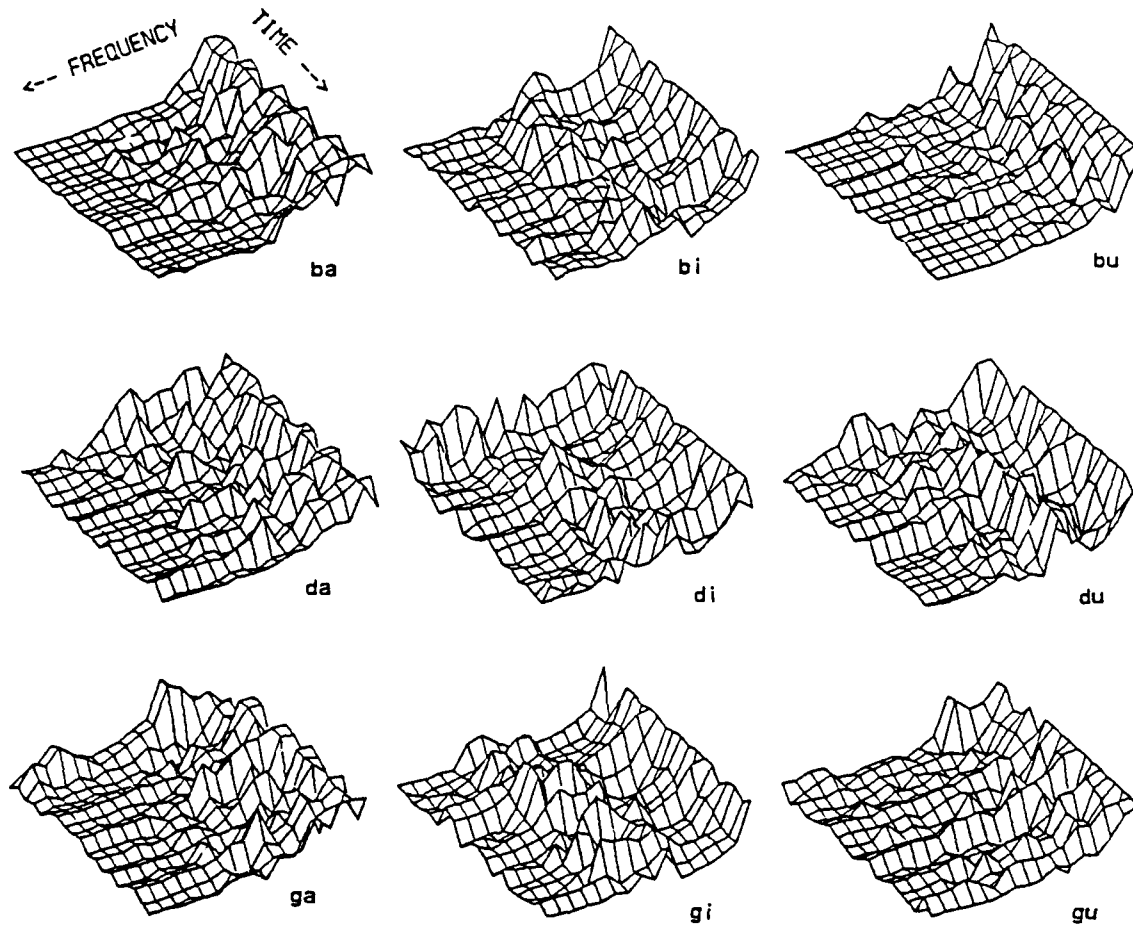


FIGURE 2. Graphs of examples of the nine sounds [ba], [bi], [bu], [da], [di], [du], [ga], [gi], and [gu]. To generate this data the release of each stop was located under computer control, and syllables were edited to include 5 ms before the release and sufficient time following release to allow 64 ms of FFT. A 64-point FFT was carried out over Hamming-windowed frames, advancing 32 points per frame. Each token consisted of an array of twenty 3.2 ms frames; within each frame the FFT output was reformatted to give spectral magnitudes over 16 evenly spaced frequency ranges, then normalized to the average for the token. These values were then transformed using the logistic function $y = 1 / (1 + e^{-x})$. This procedure maps all values to the range 0 to 1, and the average to about 0.46. Since the median for the amplitude spectra is generally just a bit greater than the average this procedure guarantees that about half of all values are mapped to each side of 0.5, which is beneficial for back-propagation learning.

hidden layer to the output layer; (c) the error for each output unit was calculated and back propagated through the network, and the weights were adjusted according to the learning rule.

The trained network was tested by scanning through all members of the training and naive data sets. The result for a token was scored as correct only if the value of the output unit that should be 1 was greater than 0.5 and the values of all the other output units were less than 0.5.

Three versions of simple phonetic labeling were run. In one, nine labels were used, corresponding to the nine sound types. In the other two, three labels were used, corresponding either to the three vowels (ignoring the consonant in the same syllable) or else to the three consonants (ignoring the vowels). In all cases more than 100,000 training cycles were run. This extensive training was used because we were interested in ultimate performance.

In all three versions, the networks always were able to learn to label the training set perfectly; that is, there were no errors in the classification of either sound type (the syllable, the vowel, or the consonant). When presented with the data from the naive set, the network trained to label whole syllables made an

average of 16% errors. The networks trained to label vowels and consonants made an average of 1.5% and 7.9% errors, respectively. The results for the vowels and consonants were about the same whether two or three hidden units were used.

We found that the performance could be improved by introducing certain kinds of noise into the samples. Simply adding random noise to the inputs degraded performance. However, if the samples were distorted by adding noise proportional to each value, performance was significantly improved. This random distortion was accomplished by replacing each input value x by $(x + xr)$, where r is randomly chosen from the range -0.5 to 0.5. When tested without noise the training sets still learned perfectly. On naive data, the performance for syllables, vowels, and consonants was 10%, 0.3% and 5.0% errors, respectively. This means that a recognizer based on vowels and consonants would have an accuracy of about 95%.

There are several reasons why training in noise should increase the model's ability to generalize. First, the noise effectively expands the data set; each syllable is represented by a larger number of exemplars. Second, and probably more important, the introduction of noise helps to blur stimulus ideosyncracies that might be learned in place of the phonetically valid features. This results in greater error during the teaching phase, but better generalization.

Now let us turn to the hidden units. They restructure the input patterns in such a way as to provide input for the final (output) layer. In the process of carrying out this mapping, they encode the input patterns as feature types. One can ask what sorts of features become represented in the hidden units as a result of the teaching phase. These internal representations may provide a clue as to how the phonetic categorization is accomplished.

In order to visualize the relationship between hidden unit activity and input sound type we used a technique that displays the average activity of each hidden unit at a different spatial position for each sound type. Examples of the hidden unit activity patterns obtained in this way are shown in Figure 3. Every hidden unit has become absolutely correlated with a subset of sound types. Hidden units have outputs of 1 for some sound types in this set and 0 for others. In addition some hidden units produce a wide range of output values for tokens not in the absolute correlation set. The correlation subsets can be vowel-like or consonant-like, in that a unit is completely on or completely off for some consonant or vowel. In the example illustrated for the nine label case, for example, two of the hidden units are vowel-like and two consonant-like. Not infrequently a unit cleanly represents a single vowel or consonant in its on activity. It is interesting that each time the learning procedure is rerun, using different random initial weights, a different pattern of hidden unit correlations is observed. However, while several unit patterns occur often, some never appear at all. For example, no hidden unit has ever been found that represents the [u] sound alone by an on unit. This contrasts with [a] and [i] which can be so represented.

Another version of label learning was carried out in which a larger number of phonetic labels was employed, reflecting a finer-grain phonetic analysis. Each of the nine syllables was divided into a *consonantal* portion and a *vocalic* portion. The *consonantal* stimulus corresponded to the first 32 ms of the syllable (starting 5 ms before release of closure) and the *vocalic* stimulus corresponded to the 32 ms of the syllable that occur 150 ms after the release of closure. This yielded 18 new stimuli. Each of the 18 stimuli types was given its own digital label, with labels randomly assigned to nine-bit codes. A network consisting of 320 input units, 6 hidden units, and 9 output units was trained on 1,000,000 learning cycles of these 1010 (505×2) stimuli.

The correlations between hidden unit activity and sound type is displayed in Figure 4. There are 6 columns, corresponding to the 6 hidden units, and 18 rows, corresponding to the 18 stimulus types. The phonetic segment is indicated by an upper case letter, and its context by a lower case letter. Thus, "Ba" refers to tokens of a voiced bilabial stop, extracted from the syllable [ba]; whereas "dI" refers to tokens of a high front vowel, extracted from the syllable [di].

Figure 4 allows us to look at the internal representation that has been developed in order to encode the 1010 tokens as 18 phonetic types. The representation is interesting in several respects. First, we see that one hidden unit (Unit 4) is always on for the first nine types, and off for the last nine types. It thus serves as a Consonant/Vowel detector. Note that the learning task has not explicitly required that

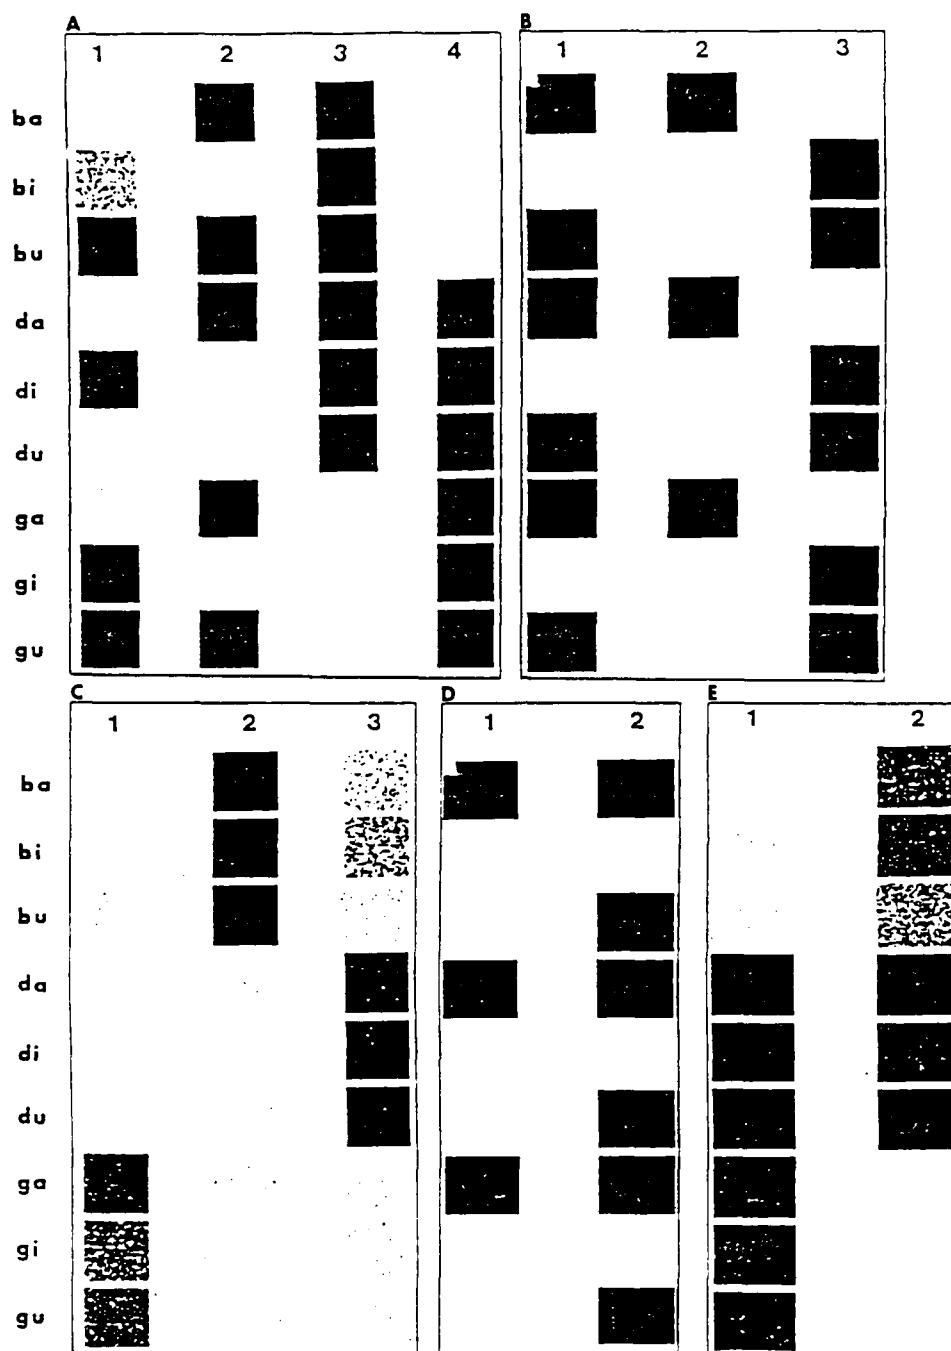


FIGURE 3. Panels A-E display hidden unit response patterns for different versions of the back propagation phonetic labeling networks. Each column shows the behavior of a single hidden unit for all nine sounds. The activity of the units is coded in the degree of darkening of the rectangle associated with each sound. A completely black rectangle indicates a unit with average activity of about 1.0 for that sound. Likewise a white rectangle (not delineated against the background) indicates an average activity near 0.0. The shaded rectangles indicate intermediate average activities. Panel A is from a four hidden unit network trained with nine labels signifying syllables. Panels B and D are from networks trained with three vowel labels, while C and E were trained to recognize three consonants.




FIGURE 4. A graph of the hidden unit activity associated with each of the 18 speech inputs. Each column shows the behavior of a single hidden unit for all 18 speech sounds. Upper case letters indicate which portion of a CV syllable was presented (consonant or vowel); lower case letters indicate the context.

this distinction be drawn. The network has simply been asked to learn 18 phonetic labels. It happens that the Consonant/Vowel is a useful dimension along which to classify types, and this dimension is implicit in the stimuli. In other cases we see that it is not as easy to interpret single hidden units by themselves. When both Unit 2 and Unit 4 are on, a velar stop (Ga, Gi, Gu) is signaled; otherwise, the vowel [i] is indicated.

One very striking result is the response pattern for Unit 0. This unit is always on (and only on) for the alveolar stops (Da, Di, Du). What makes this so surprising is that the alveolar stops exhibit a great deal of acoustic variability across different vowel contexts. The task simply required that the network learn labels for the three different alveolar allophones; it was not required to group these together in any way (indeed, there was no feedback that informed the network that any relationship existed between these three types). Nonetheless, the network has spontaneously determined a representation in which these allophones are grouped together.

The weights connecting the input to the hidden units are a kind of filter through which the sound stimuli pass to determine hidden unit activity. The shape of this filter is indicative of the sound features recognized by the hidden units. Examining these weight profiles can give us some understanding of these features. Figure 5 shows graphs of the input weights for hidden units with outputs at 1 for only a single sound or a pair of sounds. In the column on the left, the weights for several vowel-

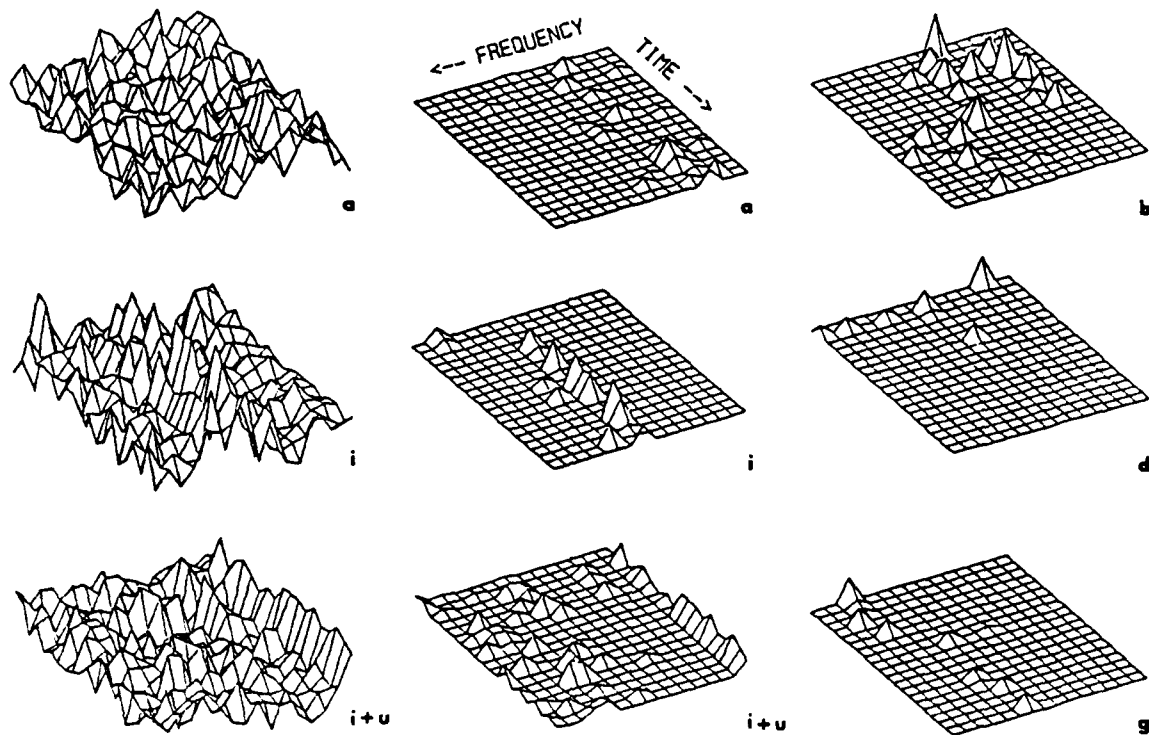


FIGURE 5. These are graphs of the weights connecting the input sound array to the hidden units. Each graph represents a single hidden unit. The weights are plotted at positions that correspond to the frequency and time of their attached input. The format of time and frequency in these graphs is the same as in Figure 2. The hidden units are all from networks trained to recognize either vowels or consonants. The hidden units are on only for the sounds indicated in the lower right-hand corner of each graph. The left hand column has complete graphs. The other two columns have flood graphs, representing only that portion of the graph above a high tide of 0.7 the whole range of weight values. All values below the high tide value are set to the high tide value. The center column is based on the same data as the one to the left. The complete graphs for the right-hand column are not shown.

recognizing units are depicted. The patterns are very complex and little can be gleaned from them. In the next column a more interpretable "flood" plot of the same data is shown; the flood plot shows only those peaks above some "high tide" level. The important differences between the weight arrays become apparent, showing some of the basis for distinguishing between the various sounds. The last column on the right is a flood plot of some consonant-recognizing units. The flood plots of the vowel-recognizing units and the consonant-recognizing units reveal only part of the story. The negative peaks are also important in the recognition process. And the importance of the finer scale structure of the weight matrix is not yet known.

Results from these studies of phonetic label learning indicate that this approach has considerable power and can be successful even given a highly confusable set of stimuli. Furthermore, the back-propagation technique results in internal representations that have interesting properties.

One important difficulty with this approach is the origin of the phonetic labels. The direct teaching technique requires that for each speech stimulus the correct label be known. It would seem desirable not to have to make this assumption. For example, from the viewpoint of child language acquisition we are put in a bit of a teleological quandry if we must assume that children know the labels of the sounds they are learning, before they learn them. This consideration led us to investigate identity mapping as a way to learn phonetic features.

IDENTITY MAPPING OF SPECTROGRAMS

In the previous study, the input and output patterns were different; the input was a known speech stimulus and the output was its abstract phonetic label. This interpretation of input and output are neither available nor necessary to the operation of the learning algorithm. It is simply learning a function that relates two patterns. One can apply the learning algorithm in a different mode in which the input and output pattern are the same. This mode, which we call *identity mapping* (it is also known as *auto-association* [Ackley, Hinton, & Sejnowski, 1985]), does not require an external teacher. Identity mapping a large pattern via a layer of a few hidden units has been shown to yield useful internal representations that give explicit information about the structure of the input patterns (Cottrell, Munro, & Zipser, in press; Zipser, in press). We have used identity mapping with the same sounds described above to see if useful internal representations can be learned in the absence of a priori knowledge about the "meaning" or "names" of patterns. We find that the hidden units in identity mapping come to represent both previously identified speech features and new, not easily described, features.

The stimuli for this experiment were identical to those used in previously. They consisted of 505 tokens of the nine consonant-vowel (CV) syllables, represented by normalized and squashed power spectra. The network had 320 input units, between 2 and 10 hidden units, and 320 output units. The training phase was similar to the labeling studies, except that the target output pattern was always identical to the input pattern.

In Figure 6A we see an example of the hidden unit activations that developed after about 150,000 learning cycles. Unit 3 is a vowel unit since it is strongly on for all [a] and off for the other vowels. Units 2 and 4 are consonant-like units since they are on quite strongly for two consonants and off for a third. Unit 3 is vowel-like but encodes some consonant information also. Units 5 and 6 cannot be characterized in terms of vowels and consonants; they represent some feature that is not easily described.

Different hidden unit activation patterns are obtained on each independent run of the same learning problem, but the same general kinds of hidden units are found. Sometimes hidden units represent a single vowel or consonant. More often they represent a strongly correlated encoding of mixed sound types as is the case with Unit 1. Units like 5 and 6, which recognize some enigmatic feature are also quite common. In general, the fewer the number of hidden units, the more strongly correlated with sound types they become.

While the identity mapping network we have described was trained on speech that was not phonetically labeled, the speech tokens had been laboriously presegmented into syllables. It occurred to us that the identity mapping network might be able to segment continuous speech. The reason this might be possible is that error would be expected to be at a minimum when the sounds used for training were in register on the input units. These error minima would then signal the boundaries between syllables.¹

To test this possibility we synthesized a pseudo-continuous speech by stringing together examples of the nine sound types in random order and shifting this sequence through the input one time-step per cycle. This resulted in an stimulus that had a complete sound token correctly in register with the inputs only once every 20 cycles. On all the rest of the cycles the input consisted of part of the end of one token and part of the beginning of another. Networks that had been fully trained to identity map presegmented sounds were used for this study, but their learning mechanism was turned off. On each shift cycle the total error was computed. (This error is just the sum-squared difference between the input and output patterns.) The results are shown in Figure 6B.

The error signal has a clear periodic component, decreasing to an identifiable minimum each time a single token is in register with the input. The reason for this is that when a CV syllable is in sync, the network "recognizes" one of the input patterns it has been trained previously on. This results in low error. On the next testing cycle, the shifted input pattern still resembles one of the learned patterns so error is relatively low, but as the shifting stimulus gets increasingly out of registration, the error

¹ Kohonen, Riittinen, and Haltsonen (1984) have used a similar scheme in their speech recognition system




FIGURE 6. A. Hidden unit response patterns from a network that had been trained to identity map the nine syllables listed on the left side. The average activity values are encoded in the same way as in Figure 3. B. Strip-chart-like plots of the total error from an identity mapping network as sound tokens are continuously shifted through the input space as described in the text. The network is the same one whose hidden unit activities are shown in A above. The "<" marks indicate times when syllables are in register in the input space. The lower panel is a continuation of the right end of the upper one.

increases. The results of this simplified study indicate that identity mapping networks could be used for segmenting continuous speech. One can also envisage using multiple networks to process speech in parallel. A network trained on identity mapping could be used to locate syllable boundaries; an error minimum could then be used to activate analysis by a second network that had been trained to do phonetic labeling.

We have seen that identity mapping can be used to learn salient features without labeling and may also be useful in segmenting speech. But we are limited by the need to presegment the input for training purposes. This is an undesirable limitation because it requires a teaching environment that may be richer than that available to the human learner. In the next section we try to remove the requirement for presegmentation of the sound stimulus.

IDENTITY MAPPING OF CONTINUOUS SPECTROGRAMS

The goal in this study was to see what sort of representation might result if a network were trained on continuous speech. As in the previous study, the task is to identity map the input. In this case, we drop the restriction that the input must correspond to a CV syllable. Instead, the speech input consisted of a corpus of 15 minutes of running speech. A text was created that contained (a) the digits from 0 to 9, (b) the 500 most frequent words from the Kucera-Francis corpus (Kucera & Francis, 1967), (c) a phonetically balanced word list of 100 items, and (d) a prose passage. The corpus was read in a conversational manner by a male speaker at a moderate rate under relatively quiet conditions, filtered at 3.5 kHz, and digitized at a 10 kHz sampling rate. The speech was hamming-windowed and analyzed by a 128-point FFT using overlapping windows that advanced 64 points per frame. The power spectra were reduced to 32 frequency bins, normalized, and squashed in a manner similar to that described above.

The network was made up of three layers; the input layer contained 640 units, the hidden layer had 8 units, and the output layer contained 640 units.

Given the magnitude of the corpus, the computational requirements of learning such a database to an acceptable level of error are considerable. Pilot studies ran for approximately 2 weeks on a VAX 11/750 digital computer (with FPA). For this reason we carried out further studies on the Cray XMP-4 computer of the San Diego Super Computer Center.

The network was trained on the corpus for 1,000,000 learning cycles. We experimented with two modes of presentation. In one mode, the speech was passed through the input layer of the network in a continuous fashion; after each learning cycle the input layer of the network was advanced by one time interval so that there was considerable overlap from one cycle to the next. In the second mode, a section of the corpus was selected at random for identity mapping; eventually all possible (overlapping) portions of the utterance were seen by the network. In pilot work, we found no differences between the two modes; the random mode is the one we adopted for our studies.

At the completion of the learning phase, the network had been trained on an extensive body of speech. The speech corpus contains approximately 140,000 different input patterns (each pattern consisting of 640 numbers). Our hope was that this sample was both representative of the variety of speech patterns for the speaker, while at the same time containing enough regularity that the network would be able to successfully encode the patterns. One graphic view of the representation that is built up is shown in Figure 7. At the bottom of the figure we see a spectrogram of a section of the training corpus. Shown above it are eight lines that graph the activations of the hidden units when the speech shown at the bottom is passed through the network. These plots can be thought of as a kind of feature representation of the speech. It is clear that features have steady-states that last for roughly syllable-sized periods of time. It has proved difficult, however, to give an interpretation of the content of these features. An important question is how much information has been preserved by the encoding contained in the eight hidden units. One can test this by seeing whether it is possible to teach a network the phonetic labels for speech sounds when we use the hidden unit representation of the identity mapped speech rather than the speech itself. This involves two steps. First, we take the nine CV syllables used before, pass them through the network previously trained on the Cray (using the 15-minute speech corpus), and then save the hidden unit activations that result. In the second step, we use the hidden unit activations to train a second network to label the activation patterns as [ba], [bi], [bu], etc. This step is analogous to the previous labeling studies, with the important difference that the input now is not speech but the representation of speech derived from the Cray-trained network. The labeling network had 560 input units (to accommodate 70 time slices, each time slice lasting 6.4 ms and being represented by eight hidden unit values), 8 hidden units, and 9 output units. The nine output units were used to encode the nine different syllable types.

After approximately 100,000 learning cycles the hidden unit activity of the labeling network had a reasonably distinct pattern that distinguished the nine different syllables. A more rigorous test is to see how many categorization errors are made by this network. There is an overall error rate of 13.5% (false




FIGURE 7. Hidden unit activity from a network that has been trained on a large, unsegmented speech corpus. The output activity of the eight hidden units is plotted on strip chart-like graphs as a sample of the speech corpus is shifted through the input space of the network. The bottom segment of the figure shows an FFT of the sound to which the hidden units are responding. The vertical timing lines are 600 ms apart. The numbers 1-8 on the left side indicate the individual hidden units.

reject = 10.4%, false accept = 3.1%), but most of it is due to a very high error rate in labeling [ga]. If this syllable is removed, the overall error rate drops to 5.8%.

This result is quite encouraging. It indicates that a degree of dimension reduction has been achieved by the corpus trained network using only eight hidden units, without an enormous loss of information. The encoded representation is rich enough that the identity of the original speech can still be extracted. Furthermore, we also see that we were able to create a feature representation without segmenting the data or knowing its phonetic identity. Indeed, features obtained in this way may eventually be useful for the task of segmentation.

One assumption we made in these studies was that the power spectra of the speech provides a good initial representation of the speech. This is not an unreasonable assumption, and there are in fact many additional assumptions one might have made (such as the use of critical bands, pre-emphasis of higher frequencies, etc.). Nonetheless, an important goal of this investigation has been to see how much of the structure of speech could be discovered with minimal a priori assumptions about what were meaningful transformations or representations of the data.

In this spirit, we wondered what might happen if we abandoned the use of the FFT to train the network. Suppose we simply presented the network with the unanalyzed digitalized waveforms? This is what we did in the final study.

IDENTITY MAPPING OF CONTINUOUS RAW SPEECH

In this study we attempted to see whether a useful input representation could be built up from continuous speech using the digital waveform itself.

The first speech corpus we used consisted of the simple sentence, *This is the voice of the neural network* (with a duration of approximately 4 seconds). The speech was kept as a series of 16-bit samples,

each value representing an A/D converter voltage, with samples occurring at 100 microsecond intervals. The network was made of 50 input units, 20 hidden units, and 50 output units. As in the previous experiment, random sections of the corpus were selected and presented as input and target for identity mapping. Since each such section contained 50 samples, the network's input window covered 5 ms of speech.

The use of pulse code modulated (PCM) speech in identity mapping makes it very easy to test the network's performance simply by collecting the output and converting it back to analog form. This requires that the output layer contain linear units; so, while the hidden layer remained nonlinear, the output layer was linear.

After one million learning cycles, we froze the weight values and fed the whole training data set through the network as a sequential stream of nonoverlapping 50 sample inputs. The output was converted to analog form and played over speakers. The result was high-quality speech; spectrograms of both the input and output are shown in Figure 8.

We were interested in seeing how well the network weights would generalize to novel speech. To test this, we retrained the network using 4 minutes of the full speech data base that was used in the Cray training study (but in PCM form, rather than as power spectra). We reasoned that this larger and more varied training set would be needed in order to learn features that would have general applicability. Learning proceeded for one million cycles, using the same presentation method as with the simple sentence. It is worth noting that because this data set contained 2.5 million different 50-sample patterns less than half the data was seen, and any pattern that was presented was typically seen only once.

The resultant network was then used as a filter for the original neural network sentence. A spectrogram of the output is shown in Figure 9; it is somewhat degraded compared with the filter that is trained on the sentence itself, but is still quite understandable.

One natural question to ask is what kind of encoding the hidden units have discovered. In Figure 10 we see spectrograms of the outputs of the individual units in response to the neural network sentence. It is clear that the responses do not resemble single sines or cosines, showing that the units have not learned a Fourier decomposition. One thing that is striking is the extent to which the hidden units' spectral responses are similar. If one compares the spectrogram of the input sentence itself (Figure 8) with those of the hidden units, one sees the way in which this is so. Most of the hidden unit response frequencies tend to center around the regions of the spectrum that are relevant for speech (this is not the case for all units; there are some that are distinctly different). The units have thus concentrated mainly on those areas of the spectrum that are relevant for encoding the speech data. The results obtained here




FIGURE 8. Spectrograms of both the input and output a network trained to identity map PCM speech.




FIGURE 9. A spectrogram of the output of the network of Figure 8 given as input the "neural network" sentence that it had never seen.

with sound are analogous to those found for visual images by Cottrell, Munro, and Zipser (in press). Since these authors were able to demonstrate considerable bandwidth compression using the hidden unit representations, we would expect that bandwidth compression has also occurred for sound.

The time scale of the hidden unit features is quite short (the features encode events on the order of 5 ms). We are primarily interested in features, but also in representations that encode larger events. Therefore, we constructed a second level network that took as its input the hidden unit representation derived from the first network and identity mapped it to an output layer. This second network had 400 input units, 10 hidden units, and 400 output units. It was able to look at representations corresponding to 100 ms (i.e., it saw 20 groups of 20 hidden unit inputs at a time, and each group of 20 hidden units came from 5 ms of speech in the first network).

The first network in this two-level system was the one trained on the extended speech corpus, and had its weights fixed. It was given a repeated sequence of nonsense syllables, [ba] [ba] [ba] [ba] (in PCM form), as a continuous stream of nonoverlapping 5 ms inputs. Every 5 ms the 20 hidden unit activations from this first network were passed to the input layer of the second network. After 400 such units had been collected by the second network, it did one cycle of learning in an identity mapping mode. On each succeeding learning cycle the input pattern in the second network was shifted left by 20 units, and a new 20 units were received from the first network. This sequence of events continued until approximately 800,000 learning cycles had occurred in the second network (remember that the first network had already been trained and was not subject to learning; it was simply acting as a pre-processor that formatted the speech in a manner analogous to the FFT used previously).

After the second network was trained on the hidden unit representations, we froze the weights in the second network. We then ran the input through both networks in a continuous stream. As we did this, we examined the pattern of 10 hidden unit activations in the second network. These hidden unit patterns indicate the syllable onsets. Our goal now is to explore the usefulness of this higher level of representation for recognizing larger sets of speech units, although at this point we remain noncommittal about what those speech units will be.

CONCLUSION

The series of experiments reported here are clearly preliminary in nature. There are a large number of questions that are raised by this work, and it is easy to think of many alternative ways of posing the problems we have presented the networks. Nonetheless, we find this approach to discovering the hidden structure in speech exciting for a number of reasons.

Power. The domain of speech processing is an extremely difficult one. There are a large number of problems of both practical and theoretical nature which remain unsolved. We believe that the experiments described here demonstrate that the PDP framework and the back-propagation method for learning are extremely powerful.

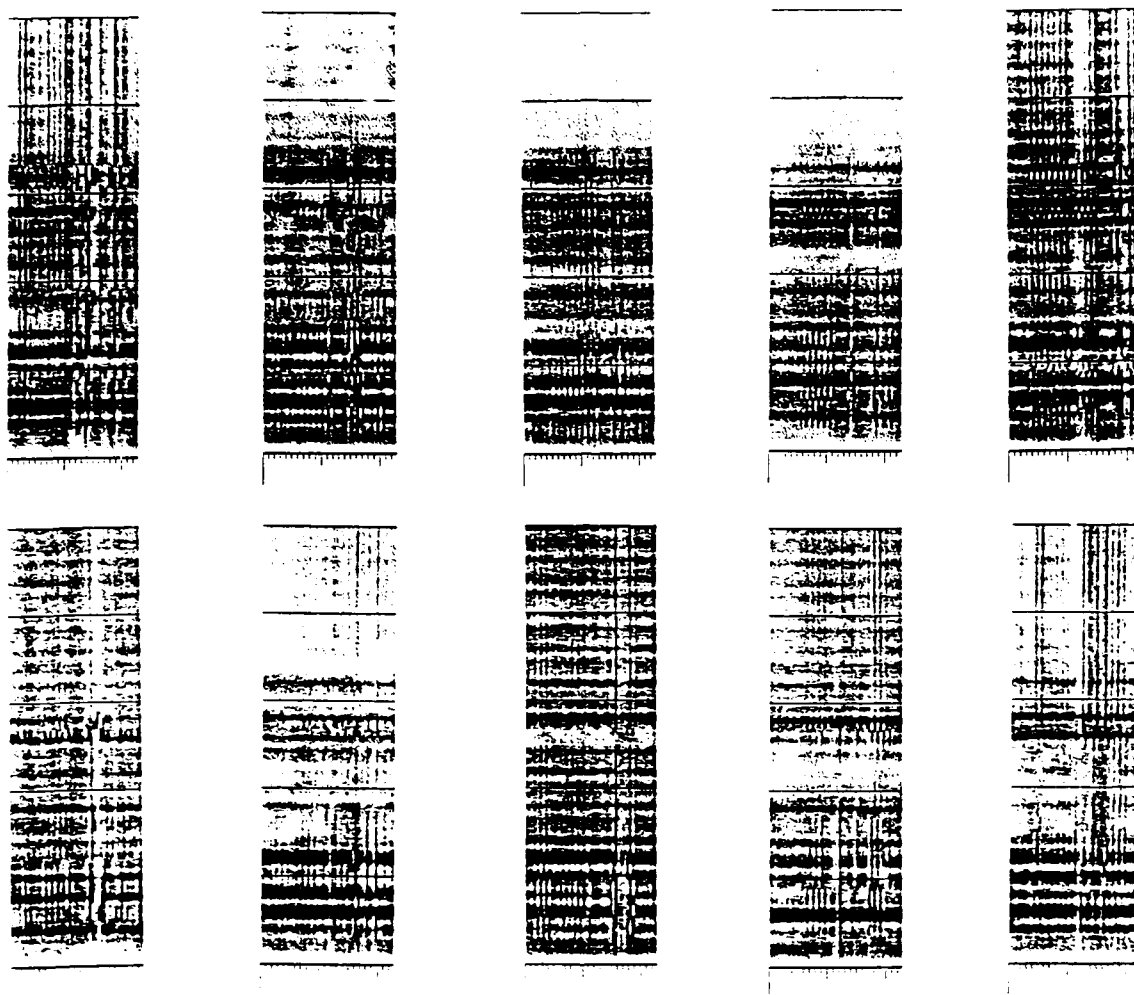


FIGURE 10. Spectrograms of the outputs of the 10 (of the 20) individual hidden units, from the network of Figure 9, in response to the "neural network" sentence.

In the first labeling study we saw that it was possible to build a system that could be taught to correctly categorize a number of highly confusable phonetic segments; and that having learned this, the network was able to generalize the categorization to novel data. We are optimistic that the performance—which was good—can be improved with refinements in the technique. We are particularly impressed with the fact that an encoding was found in which one hidden unit became active whenever an alveolar stop was presented, regardless of vocalic context, and of another which did the same for velar stops. It is well known that both of these consonants exhibit a great deal of contextual variability, and the spontaneous discovery of an invariant feature is surprising and gratifying.

Representations. An important goal in this work was to study the representations that result from applying the back-propagation learning algorithm to speech. We feel this is an important area, which we have just begun to study. In some cases the representations that are discovered are intuitively sensible and easy to interpret. In the labeling studies the Consonant/Vowel distinction was encoded by a single unit. In other cases we saw that the representation itself assumed a distributed form, with groups of hidden units participating in (for example) the encoding of place of articulation. It is interesting that

the specific representations may vary when learning experiments are replicated. This suggests that multiple networks learning the same data may provide a richer representation than any single network. Finally, we saw in the PCM identity mapping studies that the algorithm finds solutions that are highly expedient. The spectral decomposition that was carried out was clearly tuned for speech; the majority of units had responses that focused on several regions of the spectrum, and these were precisely those regions that are highly relevant for speech.

Innate vs. learned representations. We do not believe that the work described here necessarily makes strong claims that the perceptual representations by humans are learned. On the other hand, we believe that the work does argue *against* making strong claims that such representations must be innate. The tendency, in linguistics perhaps more than psychology, has been to assume that much of the representation apparatus used in processing language must be learned. In large part, that is because it has seemed to many people that the representations are complex and often arbitrary and that the input data available to the language learner for those representations is impoverished.

We feel that this study encourages the belief that more information about the structure of speech is extractable from the input than has been supposed. The back-propagation method of learning may not in fact be what is used by humans. Still, it at least demonstrates that one relatively simple algorithm does exist that is capable of uncovering a great deal of structure in a small sample of speech. It is our hope, based on these preliminary studies, that it will be possible to construct a hierarchy of learning networks that will spontaneously learn to recognize speech using only extensive examples of input speech, loosely synchronized with transcribed text. We further hope that this task can be accomplished in such a way as to shed light on the actual mechanisms used by the brain.

REFERENCES

- Ackley, D., Hinton, G. E., & Sejnowski, T. (1985). A learning algorithm for Boltzmann machines. *Cognitive Science*, 9, 147-169.
- Aronoff, M. (1976). Word formation in generative grammar. *Linguistic Inquiry Monograph 1*. Cambridge, MA: MIT Press.
- Cole, R.A., Stern, R.M., & Lasry, M.J. (1986). Performing fine phonetic distinctions: Templates vs. features. In J. S. Perkell & D. H. Klatt (Eds.), *Invariance and variability in speech processes*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Cottrell, G. W., Munro, P. W., & Zipser, D. (in press). Image compression by back propagation: A demonstration of extensional programming. In N. E. Sharkey (Ed.), *Advances in cognitive science* (Vol. 2). Norwood, NJ: Ablex.
- Dixon, N. R., & Silverman, H. F. (1976). The 1976 modular acoustic processor (MAP). *IEEE Transactions in Acoustics, Speech, and Signal Processing*, 25, 367-378.
- Fujimura, O., & Lovins, J. B. (1978). Syllables as concatenative phonetic units. In A. Bell & J. B. Hooper (Eds.), *Syllables and segments*. Amsterdam: North-Holland.
- Klatt, D.H. (1980). Scriber and LAFS: Two new approaches to speech analysis. In W. A. Lea (Ed.), *Trends in speech recognition*. Englewood Cliffs, NJ: Prentice-Hall.
- Kohonen, T., Riittinen, H., Reuhkala, E., & Haltsonen, S. (1984). On-line recognition of spoken words

from a large vocabulary. *Information Sciences*, 33, 3-30.

Lowerre, B.T. (1976). *The Harpy speech recognition system*. Unpublished doctoral dissertation, Carnegie-Mellon University, Pittsburgh.

Mehler, J. (1981). The role of syllables in speech processing: Infant and adult data. *Philosophical transactions of the Royal Society, Series B*, 295, 333-352.

Pisoni, D.B. (1981). In defense of segmental representations in speech perception. *Journal of the Acoustical Society of America*, 69 (S1), S32.

Rosenblatt, F. (1962). *Principles of neurodynamics*. New York: Spartan.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. 1. Foundations* (pp. 318-362). Cambridge, MA: MIT Press/Bradford Books.

Wickelgren, W. A. (1969). Context-sensitive coding, associative memory, and serial order in (speech) behavior. *Psychological Review*, 76, 1-15.

Zipser, D. (in press). Programming neural nets to do spatial computations. In N. E. Sharkey (Ed.), *Advances in cognitive science* (Vol. 2). Norwood, NJ: Ablex.

ICS Technical Report List

The following is a list of publications by people in the Institute for Cognitive Science. For reprints, write or call:

Institute for Cognitive Science, C-015
University of California, San Diego
La Jolla, CA 92093
(619) 534-6771

- 8301. David Zipser. *The Representation of Location*. May 1983.
- 8302. Jeffrey Elman and Jay McClelland. *Speech Perception as a Cognitive Process: The Interactive Activation Model*. April 1983. Also published in N. Lass (Ed.), *Speech and language: Volume 10*, New York: Academic Press, 1983.
- 8303. Ron Williams. *Unit Activation Rules for Cognitive Networks*. November 1983.
- 8304. David Zipser. *The Representation of Maps*. November 1983.
- 8305. The HMI Project. *User Centered System Design: Part I, Papers for the CHI '83 Conference on Human Factors in Computer Systems*. November 1983. Also published in A. Janda (Ed.), *Proceedings of the CHI '83 Conference on Human Factors in Computing Systems*. New York: ACM, 1983.
- 8306. Paul Smolensky. *Harmony Theory: A Mathematical Framework for Stochastic Parallel Processing*. December 1983. Also published in *Proceedings of the National Conference on Artificial Intelligence, AAAI-83*, Washington DC, 1983.
- 8401. Stephen W. Draper and Donald A. Norman. *Software Engineering for User Interfaces*. January 1984. Also published in *Proceedings of the Seventh International Conference on Software Engineering*, Orlando, FL, 1984.
- 8402. The UCSD HMI Project. *User Centered System Design: Part II, Collected Papers*. March 1984. Also published individually as follows: Norman, D.A. (1984), Stages and levels in human-machine interaction, *International Journal of Man-Machine Studies*, 21, 365-375; Draper, S.W., The nature of expertise in UNIX; Owen, D., Users in the real world; O'Malley, C., Draper, S.W., & Riley, M., Constructive interaction: A method for studying user-computer-user interaction; Smolensky, P., Monty, M.L., & Conway, E., Formalizing task descriptions for command specification and documentation; Bannon, L.J., & O'Malley, C., Problems in evaluation of human-computer interfaces: A case study; Riley, M., & O'Malley, C., Planning nets: A framework for analyzing user-computer interactions; all published in B. Shackel (Ed.), *INTERACT '84, First Conference on Human-Computer Interaction*, Amsterdam: North-Holland.

1984; Norman, D.A., & Draper, S.W., Software engineering for user interfaces, *Proceedings of the Seventh International Conference on Software Engineering*, Orlando, FL, 1984.

8403. Steven L. Greenspan and Eric M. Segal. *Reference Comprehension: A Topic-Comment Analysis of Sentence-Picture Verification*. April 1984. Also published in *Cognitive Psychology*, 16, 556-606, 1984.
8404. Paul Smolensky and Mary S. Riley. *Harmony Theory: Problem Solving, Parallel Cognitive Models, and Thermal Physics*. April 1984. The first two papers are published in *Proceedings of the Sixth Annual Meeting of the Cognitive Science Society*, Boulder, CO, 1984.
8405. David Zipser. *A Computational Model of Hippocampus Place-Fields*. April 1984.
8406. Michael C. Mozer. *Inductive Information Retrieval Using Parallel Distributed Computation*. May 1984.
8407. David E. Rumelhart and David Zipser. *Feature Discovery by Competitive Learning*. July 1984. Also published in *Cognitive Science*, 9, 75-112, 1985.
8408. David Zipser. *A Theoretical Model of Hippocampal Learning During Classical Conditioning*. December 1984.
8501. Ronald J. Williams. *Feature Discovery Through Error-Correction Learning*. May 1985.
8502. Ronald J. Williams. *Inference of Spatial Relations by Self-Organizing Networks*. May 1985.
8503. Edwin L. Hutchins, James D. Hollan, and Donald A. Norman. *Direct Manipulation Interfaces*. May 1985. Also published in D. A. Norman & S. W. Draper (Eds.), *User Centered System Design: New Perspectives on Human-Computer Interaction*, 1986, Hillsdale, NJ: Erlbaum.
8504. Mary S. Riley. *User Understanding*. May 1985. Also published in D. A. Norman & S. W. Draper (Eds.), *User Centered System Design: New Perspectives on Human-Computer Interaction*, 1986, Hillsdale, NJ: Erlbaum.
8505. Liam J. Bannon. *Extending the Design Boundaries of Human-Computer Interaction*. May 1985.
8506. David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. *Learning Internal Representations by Error Propagation*. September 1985. Also published in D. E. Rumelhart, J. L. McClelland, & the PDP Research Group, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Vol. 1. Foundations*, 1986, Cambridge, MA: Bradford Books/MIT Press.
8507. David E. Rumelhart and James L. McClelland. *On Learning the Past Tense of English Verbs*. October 1985. Also published in J. L. McClelland, D. E. Rumelhart, & the PDP Research Group, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Vol. 2. Psychological and Biological Models*, 1986, Cambridge, MA: MIT Press/Bradford Books.

8601. David Navon and Jeff Miller. *The Role of Outcome Conflict in Dual-Task Interference*. January 1986.
8602. David E. Rumelhart and James L. McClelland. *PDP Models and General Issues in Cognitive Science*. April 1986. Also published in D. E. Rumelhart, J. L. McClelland, & the PDP Research Group, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. Vol. 1: *Foundations*, 1986, Cambridge, MA: MIT Press/Bradford Books.
8603. James D. Hollan, Edwin L. Hutchins, Timothy P. McCandless, Mark Rosenstein, and Louis Weitzman. *Graphical Interfaces for Simulation*. May 1986. To be published in W. B. Rouse (Ed.), *Advances in Man-Machine Systems* (Vol. 3). Greenwich, CT: Jai Press.
8604. Michael I. Jordan. *Serial Order: A Parallel Distributed Processing Approach*. May 1986.
8605. Ronald J. Williams. *Reinforcement Learning in Connectionist Networks: A Mathematical Analysis*. June 1986.
8606. David Navon. *Visibility or Disability: Notes on Attention*. June 1986.
8607. William Appelbe, Donald Coleman, Allyn Fratkin, James Hutchison, and Walter J. Savitch. *Porting UNIX to a Network of Diskless Micros*. June 1986.
8608. David Zipser. *Programming Neural Nets to Do Spatial Computations*. June 1986.
8609. Louis Weitzman. *Designer: A Knowledge-Based Graphic Design Assistant*. July 1986.
8610. Michael C. Mozer. *RAMBOT: A Connectionist Expert System That Learns by Example*. August 1986.
8611. Michael C. Mozer. *Early Parallel Processing in Reading: A Connectionist Approach*. December 1986.
8701. Jeffrey L. Elman and David Zipser. *Learning the Hidden Structure of Speech*. February 1987.

Earlier Reports by People in the Cognitive Science Lab

The following is a list of publications by people in the Cognitive Science Lab and the Institute for Cognitive Science. For reprints, write or call:

Institute for Cognitive Science, C-015
University of California, San Diego
La Jolla, CA 92093
(619) 452-6771

- ONR-8001. Donald R. Gentner, Jonathan Grudin, and Eileen Conway. *Finger Movements in Transcription Typing*. May 1980.
- ONR-8002. James L. McClelland and David E. Rumelhart. *An Interactive Activation Model of the Effect of Context in Perception: Part I*. May 1980. Also published in *Psychological Review*, 88.5, pp. 375-401, 1981.
- ONR-8003. David E. Rumelhart and James L. McClelland. *An Interactive Activation Model of the Effect of Context in Perception: Part II*. July 1980. Also published in *Psychological Review*, 89, 1, pp. 60-94, 1982.
- ONR-8004. Donald A. Norman. *Errors in Human Performance*. August 1980.
- ONR-8005. David E. Rumelhart and Donald A. Norman. *Analogical Processes in Learning*. September 1980. Also published in J. R. Anderson (Ed.), *Cognitive skills and their acquisition*. Hillsdale, NJ: Erlbaum, 1981.
- ONR-8006. Donald A. Norman and Tim Shallice. *Attention to Action: Willed and Automatic Control of Behavior*. December 1980.
- ONR-8101. David E. Rumelhart. *Understanding Understanding*. January 1981.
- ONR-8102. David E. Rumelhart and Donald A. Norman. *Simulating a Skilled Typist: A Study of Skilled Cognitive-Motor Performance*. May 1981. Also published in *Cognitive Science*, 6, pp. 1-36, 1982.
- ONR-8103. Donald R. Gentner. *Skilled Finger Movements in Typing*. July 1981.
- ONR-8104. Michael I. Jordan. *The Timing of Endpoints in Movement*. November 1981.
- ONR-8105. Gary Perlman. *Two Papers in Cognitive Engineering: The Design of an Interface to a Programming System and MENUNIX: A Menu-Based Interface to UNIX (User Manual)*. November 1981. Also published in *Proceedings of the 1982 USENIX Conference*, San Diego, CA, 1982.
- ONR-8106. Donald A. Norman and Diane Fisher. *Why Alphabetic Keyboards Are Not Easy to Use: Keyboard Layout Doesn't Much Matter*. November 1981. Also published in *Human Factors*, 24, pp. 509-515, 1982.
- ONR-8107. Donald R. Gentner. *Evidence Against a Central Control Model of Timing in Typing*. December 1981. Also published in *Journal of Experimental Psychology: Human Perception and Performance*, 8, pp. 793-810, 1982.

- ONR-8201. Jonathan T. Grudin and Serge Larochelle. *Digraph Frequency Effects in Skilled Typing*. February 1982.
- ONR-8202. Jonathan T. Grudin. *Central Control of Timing in Skilled Typing*. February 1982.
- ONR-8203. Amy Geoffroy and Donald A. Norman. *Ease of Tapping the Fingers in a Sequence Depends on the Mental Encoding*. March 1982.
- ONR-8204. LNR Research Group. *Studies of Typing from the LNR Research Group: The role of context, differences in skill level, errors, hand movements, and a computer simulation*. May 1982. Also published in W. E. Cooper (Ed.), *Cognitive aspects of skilled typewriting*. New York: Springer-Verlag, 1983.
- ONR-8205. Donald A. Norman. *Five Papers on Human-Machine Interaction*. May 1982. Also published individually as follows: Some observations on mental models, in D. Gentner and A. Stevens (Eds.), *Mental models*, Hillsdale, NJ: Erlbaum, 1983; A psychologist views human processing: Human errors and other phenomena suggest processing mechanisms, in *Proceedings of the International Joint Conference on Artificial Intelligence*, Vancouver, 1981; Steps toward a cognitive engineering: Design rules based on analyses of human error, in *Proceedings of the Conference on Human Factors in Computer Systems*, Gaithersburg, MD, 1982; The trouble with UNIX, in *Datamation*, 27,12, November 1981, pp. 139-150; The trouble with networks, in *Datamation*, January 1982, pp. 188-192.
- ONR-8206. Naomi Miyake. *Constructive Interaction*. June 1982.
- ONR-8207. Donald R. Gentner. *The Development of Typewriting Skill*. September 1982. Also published as Acquisition of typewriting skill, in *Acta Psychologica*, 54, pp. 233-248, 1983.
- ONR-8208. Gary Perlman. *Natural Artificial Languages: Low-Level Processes*. December 1982. Also published in *The International Journal of Man-Machine Studies*, 20, pp. 373-419, 1984.
- ONR-8301. Michael C. Mozer. *Letter Migration in Word Perception*. April 1983. Also published in *Journal of Experimental Psychology: Human Perception and Performance*, 9, 4, pp. 531-546, 1983.
- ONR-8302. David E. Rumelhart and Donald A. Norman. *Representation in Memory*. June 1983. To appear in R. C. Atkinson, G. Lindzey, & R. D. Luce (Eds.), *Handbook of experimental psychology*. New York: Wiley (in press).

Dr. Phillip L. Ackerman University of Minnesota Department of Psychology Minneapolis, MN 55455	Dr. Nancy S. Anderson Department of Psychology University of Maryland College Park, MD 20742	ICDR Robert Carter Office of the Chief of Naval Operations OP-01B Pentagon Washington, DC 20350-2000
Dr. Beth Adelson Department of Computer Science Tufts University Medford, MA 02155	Technical Director, ARI 5001 Eisenhower Avenue Alexandria, VA 22333	Dr. Arthur S. Blalwes Code N/11 Naval Training Systems Center Orlando, FL 32813
Air Force Human Resources Lab AFHRL/MPD Brooks AFB, TX 78235	Dr. Gary Aston-Jones Department of Biology New York University 1009 Main Bldg Washington Square New York, NY 10003	Chair, Department of Psychology College of Arts and Sciences Catholic University of America Washington, DC 20064
AFOSR, Life Sciences Directorate Bolling Air Force Base Washington, DC 20332	Dr. Alan Baddeley Medical Research Council Applied Psychology Unit 15 Chaucer Road Cambridge CB2 2EF ENGLAND	Dr. David E. Clement Department of Psychology University of South Carolina Columbia, SC 29208
Technical Director, Army Human Engineering Lab Aberdeen Proving Ground MD 21005	Prof. dott. Bruno G. Bara Unità di ricerca di Intelligenza artificiale Università di Milano 20122 Milano - Via F. Sforza 23 ITALY	Dr. Charles Clifton Tobin Hall Department of Psychology University of Massachusetts Amherst, MA 01003
Dr. Robert Ahlers Code N/11 Human Factors Laboratory Naval Training Systems Center Orlando, FL 32813	Dr. Jackson Beatty Department of Psychology University of California Los Angeles, CA 90024	Dr. Michael Coles University of Illinois Department of Psychology Champaign, IL 61820
Dr. Ed Alken Navy Personnel R&D Center San Diego, CA 92152-6800	Dr. A. L. Becker University of Michigan 1091 Fricke/Linguistics Ann Arbor, MI 48109	Dr. Stanley Collier Office of Naval Technology Code 222 800 N. Quincy Street Arlington, VA 22217-5000
Dr. John Allen Department of Psychology George Mason University 4400 University Drive Fairfax, VA 22030	Dr. Alvin Blittner Naval Biodynamics Laboratory New Orleans, LA 70189	Dr. Leon Cooper Brown University Center for Neural Science Providence, RI 02912
Dr. James Anderson Brown University Center for Neural Science Providence, RI 02912	Dr. John Black Teachers College Columbia University 525 West 121st Street New York, NY 10027	Brian Dallman 3400 TWM/TTGXS Lowry AFB, CO 80230-5000
Dr. John R. Anderson Department of Psychology Carnegie-Mellon University Pittsburgh, PA 15213	Dr. Pat Carpenter Carnegie Mellon University Department of Psychology Pittsburgh, PA 15213	Dr. Joel Davis Office of Naval Research Code 1142B1 800 North Quincy Street Arlington, VA 22217-5000
	Dr. Gall Carpenter Northeastern University Department of Mathematics, 5041A 360 Huntington Avenue Boston, MA 02115	
	Dr. John Bruer The James S. McDonnell Foundation University Club Tower, Suite 1610 1034 South Brentwood Blvd. St. Louis, MO 63117	
	Joanne Capper Center for Research into Practice 1718 Connecticut Ave., N.W. Washington, DC 20009	
	Dr. Robert Breaux Code N-095H Naval Training Systems Center Orlando, FL 32813	
	Commanding Officer CAPT Lorin W. Brown NROTC Unit Illinois Institute of Technology 3300 S. Federal Street Chicago, IL 60616-3793	
	Dr. John Bruer The James S. McDonnell Foundation University Club Tower, Suite 1610 1034 South Brentwood Blvd. St. Louis, MO 63117	

Dr. Laura Davis
NRL/NCARAI
Code 7510
4555 Overlook Ave., SW
Washington, DC 20375-5000

Dr. William Epstein
University of Wisconsin
W. J. Brodgen Psychology Bldg.
1202 W. Johnson Street
Madison, WI 53706

ERIC Facility-Acquisitions
4833 Rugby Avenue
Bethesda, MD 20014

Dr. K. Anders Ericsson
University of Colorado
Department of Psychology
Boulder, CO 80309

Dr. W. E. Evans
Hubbs Sea World Institute
1720 S. Shores Rd.
Mission Bay
San Diego, CA 92109

Dr. Martha Farah
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Marshall J. Farr
Farr-Sight Co.
2520 North Vernon Street
Arlington, VA 22207

Dr. Jerome A. Feldman
University of Rochester
Computer Science Department
Rochester, NY 14627

Dr. Gail R. Fleischaker
Marquits Lab
Biological Sci. Center
2 Cunningham Street
Boston, MA 02215

J. D. Fletcher
9931 Corsica Street
Vienna VA 22180

Dr. Jane M. Flinn
Department of Psychology
George Mason University
4400 University Drive
Fairfax, VA 22030

Dr. Gerald F. DeJong
Artificial Intelligence Group
Coordinated Science Laboratory
University of Illinois
Urbana, IL 61801

Dr. R. K. Dismukes
Associate Director for Life Sciences
AFOSR
Bolling AFB
Washington, DC 20332

Dr. Stephanie Doan
Code 6021
Naval Air Development Center
Warminster, PA 18974-5000

Dr. Emanuel Donchin
University of Illinois
Department of Psychology
Champaign, IL 61820

Defense Technical
Information Center
Cameron Station, Bldg 5
Alexandria, VA 22314
Attn: TC
(12 Copies)

Dr. Ford Ebner
Brown University
Anatomy Department
Medical School
Providence, RI 02912

Dr. Jeffrey Elman
University of California,
San Diego
Department of Linguistics, C-008
La Jolla, CA 92093

Dr. John R. Frederiksen
Bolt, Beranek & Newman
50 Moulton Street
Cambridge, MA 02138

Dr. Michael Friendly
Psychology Department
York University
Toronto ONT
CANADA M3J 1P3

Julie A. Gadsden
Information Technology
Applications Division
Admiralty Research Establishment
Portsmouth, Portsmouth PO6 4AA
UNITED KINGDOM

Dr. Michaela Gallagher
University of North Carolina
Department of Psychology
Chapel Hill, NC 27514

Dr. Michael Genesereth
Stanford University
Computer Science Department
Stanford, CA 94305

Dr. Don Gentner
Center for Human
Information Processing
University of California
La Jolla, CA 92093

Dr. Claude Ghez
Center for Neurobiology and
Behavior
722 W. 168th Street
New York, NY 10032

Dr. Lee Giles
AFOSR
Bolling AFB
Washington, DC 20332

Dr. Eugene E. Gloye
Office of Naval Research
Detachment
1030 E. Green Street
Pasadena, CA 91106-2485

Dr. Sam Glucksberg
Department of Psychology
Princeton University
Princeton, NJ 08540

Dr. Daniel Gopher
Industrial Engineering
& Management
TECHNION
Haifa 32000
ISRAEL

Dr. Sherrie Gott
AFHRL/MDJ
Brooks AFB, TX 78235

Jordan Grafman, Ph.D.
2021 Lyttonville Road
Silver Spring, MD 20910

Dr. Richard H. Granger
Department of Computer Science
University of California, Irvine
Irvine, CA 92717

Dr. Steven Grant
Department of Biology
New York University
1009 Main Bldg
Washington Square
New York, NY 10003

Dr. Wayne Gray
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. William Greenough
University of Illinois
Department of Psychology
Champaign, IL 61820

Dr. Dik Gregory
Behavioral Sciences Division
Admiralty Research
Establishment
Teddington
Middlesex, ENGLAND

Dr. Stephen Grossberg
Center for Adaptive Systems
Room 244
111 Cunningham Street
Boston University
Boston, MA 02215

Dr. Muhammad K. Habib
University of North Carolina
Department of Biostatistics
Chapel Hill, NC 27514

Dr. Henry M. Haliff
Halff Resources, Inc.
4918 33rd Road, North
Arlington, VA 22207

Dr. Ray Hannapel
Scientific and Engineering
Personnel and Education
National Science Foundation
Washington, DC 20550

Stevan Harnad
Editor, The Behavioral and
Brain Sciences
20 Nassau Street, Suite 240
Princeton, NJ 08540

Dr. Steven A. Hilliard
Department of Neurosciences
University of California,
San Diego
La Jolla, CA 92093

Dr. Geoffrey Hinton
Carnegie-Mellon University
Computer Science Department
Pittsburgh, PA 15213

Dr. Jim Hollan
Intelligent Systems Group
Institute for
Cognitive Science (C-015)
UCSD
La Jolla, CA 92093

Dr. John Holland
University of Michigan
2313 East Engineering
Ann Arbor, MI 48109

Dr. James Howard
Dept. of Psychology
Human Performance Laboratory
Catholic University of
America
Washington, DC 20064

Dr. Earl Hunt
Department of Psychology
University of Washington
Seattle, WA 98105

Dr. Alice Isen
Department of Psychology
University of Maryland
Catonsville, MD 21228

Chair, Department of
Psychology
The Johns Hopkins University
Baltimore, MD 21218

CDR Tom Jones
ONR Code 125
800 N. Quincy Street
Arlington, VA 22217-5000

Dr. Douglas H. Jones
Thatcher Jones Associates
P.O. Box 6640
10 Trafalgar Court
Lawrenceville, NJ 08648

Dr. Marcel Just
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Daniel Kahneman
Department of Psychology
University of California
Berkeley, CA 94720

Dr. Demetrios Karis
Grumman Aerospace Corporation
MS C04-14
Bethpage, NY 11714

Dr. Milton S. Katz
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22355

Dr. Steven W. Keele
Department of Psychology
University of Oregon
Eugene, OR 97403

Dr. Wendy Kellogg
IBM T. J. Watson Research Ctr.
P.O. Box 218
Yorktown Heights, NY 10598

Dr. Scott Kelso
Haskins Laboratories,
270 Crown Street
New Haven, CT 06510

Dr. David Klahr
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Sylvan Kornblum
University of Michigan
Mental Health Research Institute
205 Washtenaw Place
Ann Arbor, MI 48109

Dr. Stephen Kosslyn
Harvard University
1236 William James Hall
33 Kirkland St.
Cambridge, MA 02138

Dr. David R. Lambert
Naval Ocean Systems Center
Code 441T
271 Catalina Boulevard
San Diego, CA 92152-6800

Dr. Pat Langley
University of California
Department of Information
and Computer Science
Irvine, CA 92717

Dr. Marcy Lansman
University of North Carolina
The L. L. Thurstone Lab.
Davie Hall 013A
Chapel Hill, NC 27514

Dr. Alan M. Lesgold
Learning Research and
Development Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Alan Leshner
Deputy Division Director
Behavioral and Neural Sciences
National Science Foundation
1800 G Street
Washington, DC 20550

Library
Naval War College
Newport, RI 02940

Library
Naval Training Systems Center
Orlando, FL 32813

Dr. Gary Lynch
University of California
Center for the Neurobiology of
Learning and Memory
Irvine, CA 92717

Dr. Don Lyon
P. O. Box 44
Higley, AZ 85236

Dr. Evans Mandus
Department of Psychology
George Mason University
4400 University Drive
Fairfax, VA 22030

Dr. Elaine Marsh
Naval Research Laboratory
Code 7510
4555 Overlook Avenue, Southwest
Washington, DC 20375-5000

Dr. James McBride
Psychological Corporation
c/o Harcourt, Brace,
Jovanovich Inc.
1250 West 6th Street
San Diego, CA 92101

Dr. Jay McClelland Department of Psychology Carnegie-Mellon University Pittsburgh, PA 15213	Dr. Randy Mumaw Program Manager Training Research Division HumRRO 1100 S. Washington Alexandria, VA 22314	Fleet Support Office, NPRDC (Code 301) San Diego, CA 92152-6800	Special Assistant for Marine Corps Matters, ONR Code 00MC 800 N. Quincy St. Arlington, VA 22217-5000
Dr. James L. McLaughlin Center for the Neurobiology of Learning and Memory University of California, Irvine Irvine, CA 92717	Chair, Department of Computer Science U.S. Naval Academy Annapolis, MD 21402	Library, NPRDC Code P201L San Diego, CA 92152-6800	Psychologist Office of Naval Research Liaison Office, Far East APO San Francisco, CA 96503
Dr. Joe McLachlan Navy Personnel R&D Center San Diego, CA 92152-6800	Chair, Department of Systems Engineering U.S. Naval Academy Annapolis, MD 21402	Dr. Michael Oberlin Naval Training Systems Center Code 711 Orlando, FL 32813-7100	Dr. Judith Orasanu Army Research Institute 5001 Eisenhower Avenue Alexandria, VA 22333
Dr. James S. McMichael Navy Personnel Research and Development Center San Diego, CA 92152	Dr. David Navon Institute for Cognitive Science University of California La Jolla, CA 92093	Dr. Stellan Ohlsson Learning R & D Center University of Pittsburgh 3939 O'Hara Street Pittsburgh, PA 15213	Dr. Robert F. Pasnak Department of Psychology George Mason University 4400 University Drive Fairfax, VA 22030
Dr. Douglas L. Medin Department of Psychology University of Illinois 603 E. Daniel Street Champaign, IL 61820	Dr. Allen Newell Department of Psychology Carnegie-Mellon University Schenley Park Pittsburgh, PA 15213	Dr. James B. Olsen Director, Waterford Testing Center 1681 West 820 North Provo, UT 84601	Daira Paulson Code 52 - Training Systems Navy Personnel R&D Center San Diego, CA 92152-6800
Dr. Jose Mestre Department of Physics Hasbrouck Laboratory University of Massachusetts Amherst, MA 01003	Dr. Mary Jo Missen University of Minnesota N218 Elliott Hall Minneapolis, MN 55455	Office of Naval Research, Code 1133 800 N. Quincy Street Arlington, VA 22217-5000	Dr. James M. Pellegrino University of California, Santa Barbara Department of Psychology Santa Barbara, CA 93106
Dr. Al Meyrowitz Office of Naval Research Code 1133 800 N. Quincy Arlington, VA 22217-5000	Dr. Donald A. Norman Institute for Cognitive Science C-015 University of California, San Diego La Jolla, California 92093	Office of Naval Research, Code 1142PS 800 N. Quincy Street Arlington, VA 22217-5000	Department of Computer Science, Naval Postgraduate School Monterey, CA 93940
Dr. George A. Miller Department of Psychology Green Hall Princeton University Princeton, NJ 08540	Director, Training Laboratory, NPRDC (Code 05) San Diego, CA 92152-6800	Office of Naval Research, Code 1142CS 800 N. Quincy Street Arlington, VA 22217-5000 (6 Copies)	Dr. Steven Pinker Department of Psychology E10-018 M.I.T. Cambridge, MA 02139
Dr. Tom Moran Xerox PARC 3333 Coyote Hill Road Palo Alto, CA 94304	Director, Manpower and Personnel Laboratory, NPRDC (Code 06) San Diego, CA 92152-6800	Psychologist Office of Naval Research Branch Office, London Box 39 FPO New York, NY 09510	Dr. Martha Polson Department of Psychology Campus Box 346 University of Colorado Boulder, CO 80309
	Director, Human Factors & Organizational Systems Lab, NPRDC (Code 07) San Diego, CA 92152-6800		Dr. Peter Polson University of Colorado Department of Psychology Boulder, CO 80309

Dr. Michael I. Posner
Department of Neurology
Washington University
Medical School
St. Louis, MO 63110

Dr. Mary C. Potter
Department of Psychology
MIT (E-10-032)
Cambridge, MA 02139

Dr. Karl Pribram
Stanford University
Department of Psychology
Bldg. 4201 -- Jordan Hall
Stanford, CA 94305

Dr. Lynne Rader
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. James A. Reggia
University of Maryland
School of Medicine
Department of Neurology
22 South Greene Street
Baltimore, MD 21201

Dr. Daniel Reissberg
Department of Psychology
New School for Social Research
65 Fifth Avenue
New York, NY 10003

Dr. Gil Ricard
Mail Stop C04-14
Grumman Aerospace Corp.
Bethpage, NY 11714

Dr. David Rumelhart
Center for Human
Information Processing
Univ. of California
La Jolla, CA 92093

Dr. E. I. Saltzman
Haskins Laboratories
270 Crown Street
New Haven, CT 06510

Dr. Arthur Samuel
Yale University
Department of Psychology
Box 11A, Yale Station
New Haven, CT 06520

Dr. Walter Schneider
Leanning R&D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

Dr. Hans-Wilhelm Schrollff
Institut fuer Psychologie
der RWTH Aachen
Jaegeystrasse zwischen 17 u. 19
5100 Aachen
WEST GERMANY

Dr. Miriam Schustack
Code 51
Navy Personnel R & D Center
San Diego, CA 92152-6800

Dr. Marc Sebrechts
Department of Psychology
Wesleyan University
Middletown, CT 06475

Dr. Colleen M. Selfert
Intelligent Systems Group
Institute for
Cognitive Science (C-015)
UCSD
La Jolla, CA 92093

Dr. I. B. Sheridan
Dept. of Mechanical Engineering
MIT
Cambridge, MA 02139

Dr. Randall Shumaker
Naval Research Laboratory
Code 7310
4555 Overlook Avenue, S.W.
Washington, DC 20375-5000

Dr. Herbert A. Simon
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Gail Siemon
Logicon
P.O. Box 85158
San Diego, CA 92138

Dr. Linda B. Smith
Department of Psychology
Indiana University
Bloomington, IN 47405

Dr. Robert F. Smith
Department of Psychology
George Mason University
4400 University Drive
Fairfax, VA 22030

Dr. Kathryn T. Spoehr
Brown University
Department of Psychology
Providence, RI 02912

Dr. Ted Steinke
Dept. of Geography
University of South Carolina
Columbia, SC 29208

Dr. Saul Sternberg
University of Pennsylvania
Department of Psychology
3815 Walnut Street
Philadelphia, PA 19104

Dr. Paul J. Sticha
Senior Staff Scientist
Training Research Division
HumRRO
1100 S. Washington
Alexandria, VA 22314

Dr. Steve Suomi
NIH Bldg. 31
Room B2B-15
Bethesda, MD 20205

Dr. John Tangney
AFOSR/NL
Boiling AFB, DC 20332

Dr. Richard F. Thompson
Stanford University
Department of Psychology
Bldg. 4201 -- Jordan Hall
Stanford, CA 94305

Dr. Sharon Tkacz
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Michael T. Turvey
Haskins Laboratories
270 Crown Street
New Haven, CT 06510

Headquarters, U. S. Marine Corps
Code MP1-20
Washington, DC 20380

Dr. William Uttal
NOSC, Hawaii Lab
Box 997
Kailua, HI 96734

Dr. Paul Van den Broek
Department of Psychology
University of Kentucky
Lexington, KY 40506-0044

Dr. Kurt Van Lehn
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Jerry Vogt
Navy Personnel R&D Center
Code 51
San Diego, CA 92152-6800

Dr. Beth Warren
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

Dr. Shih-Sung Wen
Jackson State University
1325 J. R. Lynch Street
Jackson, MS 39217

Dr. Douglas Wetzel
Code 12
Navy Personnel R&D Center
San Diego, CA 92152-6800

END

4-87

DTIC